Amendments to the claims,
        Listing of all claims pursuant to 37 CFR 1.121(c)

This listing of claims will replace all prior versions, and listings, of claims in the application:

What is claimed is:

    1. (Previously presented) In a database system including a database optimizer normally used for generating an access plan for processing a given database query run against the database system, an optimizer-based method for recommending database indexes to be created for maximizing system performance, the method comprising:

        capturing a workload representative of database queries employed during prior database system use;

        monitoring the database optimizer as it prepares to optimize each of the queries, by recording all potential database physical indexes which do not currently exist in the database and for which the database optimizer searched during a preoptimization phase that occurs prior to access plan generation;

        creating an initial set of virtual indexes each simulating presence of a class of potential database physical indexes that were recorded during said monitoring step, wherein each said virtual index comprises an in-memory data structure corresponding to a set of potential database physical indexes;

        computing cost benefits for different subsets of the set of virtual indexes by invoking the database optimizer again for purposes of providing cost estimates for the workload for each such subset, progressively eliminating a fixed percentage of virtual indexes with the lowest expected improvement in the query's estimated cost from consideration until space that would be occupied by the virtual indexes is less than or equal to a user-specified value; and

        recommending database physical indexes to be created based on those virtual indexes that have favorable estimated cost benefits for the captured workload and have not been eliminated from consideration.

    2. (Canceled)

3. (Original) The method of claim 1, wherein the capturing step includes:

displaying a screen input button that a user may invoke to record a usage session as a workload.

4. (Original) The method of claim 1, wherein the workload represents user execution of a database application with a typical workload that is contemplated for the application.

5. (Original) The method of claim 1, wherein the workload includes information recording text of all the queries operating during the capture of the workload.

6. (Original) The method of claim 1, wherein the workload includes information recording settings for certain options that affect how queries are optimized.

7. (Original) The method of claim 1, wherein the capturing step includes:

capturing information about a set of workloads to define a problem instance.

8. (Original) The method of claim 1, further comprising:

setting a limit on how much disk space is available for physical indexes.

9. (Original) The method of claim 8, wherein the recommending step takes into account the limit on disk space available for physical indexes.

10. (Original) The method of claim 8, wherein the recommending step includes:

if the physical indexes to be recommended for creation exceed the limit on disk space available for physical indexes, removing some of the physical indexes from consideration.

11. (Original) The method of claim 10, wherein the physical indexes removed from consideration are ones having less favorable cost benefits for the captured workload.

12. (Original) The method of claim 10, wherein the physical indexes removed from consideration comprise at least 20 percent of bottom performing indexes considered for recommendation.

13. (Original) The method of claim 1, further comprising:
specifying whether certain types of indexes should be considered at all.

14. (Original) The method of claim 1, wherein the creating virtual indexes step includes:
searching for relevant indexes that will help the system's optimizer use sargable predicates for partial index scans.

15. (Previously presented) The method of claim 14, wherein an index consultant creates virtual indexes without specifying ordering of columns used in sargable equality predicates.

16. (Original) The method of claim 1, wherein the creating virtual indexes step includes:
searching for relevant indexes that will help provide useful orderings.

17. (Original) The method of claim 16, wherein columns of virtual indexes may be order-independent "don't care" columns that satisfy some interesting ordering wish list of the system's optimizer.

18. (Original) The method of claim 16, wherein columns of virtual indexes may have an unspecified sortedness.

19. (Original) The method of claim 1, further comprising:
collapsing some of the virtual indexes together, if feasible for the workload.

20. (Original) The method of claim 19, wherein the collapsing step includes:

identifying that columns of one virtual index are a superset of another the columns of another virtual index, and that both indexes may be combined into a single virtual index that is feasible for the workload; and

identifying that sortedness of a column of a virtual index, if unspecified, may be specified to allow it to be combined with an index with identical columns but specified sortedness; and

identifying that a virtual index that has columns of opposite sortedness of a second virtual index, and that both indexes may be combined into a single virtual index.

21. (Original) The method of claim 1, further comprising:

polling periodically in the method to ensure that the system is working with accurate cost information.

22. (Original) A computer-readable medium having processor-executable instructions for performing the method of claim 1.

23. (Original) A downloadable set of processor-executable instructions for performing the method of claim 1.

24. (Previously presented) A system that recommends database indexes to be created for optimizing system performance, the system comprising:

a database system that executes database queries, said database system including a database optimizer normally used for generating an access plan for processing each given database query run against the database system; and

an optimizer-based index consultant for capturing a workload representative of database queries executed during prior system use, which monitors the database optimizer as it prepares to optimize each of the queries by recording all potential database physical indexes that do not currently exist in the database and for which the database optimizer searched during a preoptimization phase that occurs prior to access plan generation, and which creates an initial set of virtual indexes each simulating presence of

a class of potential database physical indexes that were recorded during monitoring, wherein each said virtual index comprises an in-memory data structure corresponding to a set of potential database physical indexes, wherein cost benefits for different subsets of the set of virtual indexes are computed by invoking the database optimizer again for purposes of providing cost estimates for the workload for each such subset, progressively eliminating a fixed percentage of virtual indexes with the lowest expected improvement in the query's estimated cost from consideration until space that would be occupied by the virtual indexes is less than or equal to a user-specified value, and wherein database physical indexes are recommended to be created based on those virtual indexes that have favorable estimated cost benefits for the captured workload and have not been eliminated from consideration.

25. (Canceled)

26. (Original) The system of claim 24, wherein the index consultant displays a screen input button that a user may invoke to record a usage session as a workload.

27. (Original) The system of claim 24, wherein the workload represents user execution of a database application with a typical workload that is contemplated for the application.

28. (Original) The system of claim 24, wherein the workload includes information recording text of all the queries operating during the capture of the workload.

29. (Original) The system of claim 24, wherein the workload includes information recording settings for certain options that affect how queries are optimized.

30. (Original) The system of claim 24, wherein the index consultant captures information about a set of workloads to define a problem instance.

31. (Original) The system of claim 24, wherein the index consultant may receive

information specifying a limit on how much disk space is available for physical indexes.

32. (Original) The system of claim 31, wherein the index consultant takes into account the limit on disk space available for physical indexes.

33. (Original) The system of claim 31, wherein the index consultant removes some of the physical indexes from consideration, when sufficient disk space is unavailable.

34. (Original) The system of claim 33, wherein the physical indexes removed from consideration are ones having less favorable cost benefits for the captured workload.

35. (Original) The system of claim 33, wherein the physical indexes removed from consideration comprise at least 20 percent of bottom performing indexes considered for recommendation.

36. (Original) The system of claim 24, wherein the index consultant allows user input specifying whether certain types of indexes should be considered at all.

37. (Original) The system of claim 24, wherein the index consultant searches for relevant indexes that will help the system's optimizer use sargable predicates for partial index scans.

38. (Previously presented) The system of claim 37, wherein the index consultant creates virtual indexes without specifying ordering of columns used in sargable equality predicates.

39. (Original) The system of claim 24, wherein the index consultant searches for relevant indexes that will help provide useful interesting (order or grouping) properties.

40. (Original) The system of claim 39, wherein columns of indexes created may reflect order-independent "don't care" columns that satisfy some interesting ordering wish lists of the system's optimizer.

41. (Original) The system of claim 24, wherein the index consultant attempts to collapse some of the virtual indexes together, if feasible for the workload.

42. (Original) The system of claim 41, wherein the index consultant attempts to identify that columns of one index are a superset of the columns of another index, and that both indexes may be combined into a single index that is feasible for the workload.

43. (Original) The system of claim 24, wherein operation of the index consultant may be polled during operation to ensure that the system is working with accurate cost information.

44. (Previously presented) The method of claim 1, wherein the virtual indexes are created by an index consultant that observes the optimizer's need for certain indexes during generation of access plans for said database queries.

45. (Previously presented) The system of claim 24, wherein the index consultant creates the virtual indexes by observing the optimizer's need for certain indexes during generation of access plans for said database queries.